

review

Sseyo Koan
Generative Music Tool

by Patrick Lichty

software

Although Sseyo (now part of Tao) is focused on mobile device audio through the Tao intent Sound System (iSS), one of that company's older core technologies still has a lot of life in it: the *Koan Pro Generative Music System*. In many ways, it's like a Swiss Army knife for those who wish to do a number of things affordably; for example, create web-based generative music that can interface with Flash, Java and Javascript; or produce algorithmical-based music and sound for live and installation settings or even ringtones.

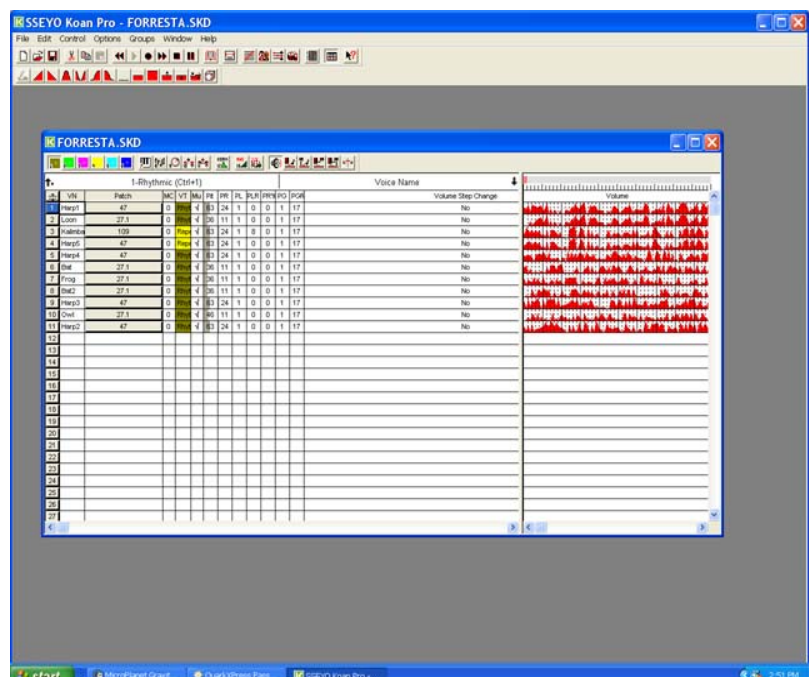
First of all, let's admit that *Koan* has a lineage that stretches back to 1991, and the latest revision (2.7) states a revision date of 2001. One can say that it's extremely robust in its functionality, very stable, not overly pretty, and a little counterintuitive, but like the cockpit of a 767. It gives you a lot of control, and the controls are obviously created by those with a technical bent rather than an interest in lyrical interfaces such as Kai Krause and Eric Wenger. Although one could argue that there are considerably more facets to the program, three main areas are of particular interest: the voice screen, the pattern editor, and the SKME (*Sseyo Koan Music Engine*). There are many subsections to each of these, such as scale, harmony, predictive determination (next note), and so on, but for the most part, I will stay with general descriptions of these three main components.

Within the Voice Editor, the main panel is comprised of the various parametric controls for the different voices, determining specific probability-based (remember that this is a generative music system) controls for the progression and shaping of the composition. These parameters include the possibility for variations in tuning / microtuning, arpeggio mutation, phrase length, scale progression sound / patch change and so on. The parameters shape six distinct types of voices: Ambient, Repeat Bar, Rhythmic, Follows, Listening and Fixed Pattern. Where the first three consist of voices that generate their own patterns, Follows and Listening are more akin to 'meat' voices that create their information from other sources, and Fixed Pattern produces compositional structures from more traditional fixed scores.

Each of the three autogenerative voices -- Ambient, Repeat Bar, and Rhythmic -- has its own distinct stylistic advantages. For slow, undulating pads, the Ambient voice allows for the creation of background swells and subtle textures. The Rhythmic voice, which is probably one of my favorites, allows for the shaping of a pattern that is the most chaotic, consisting of more amorphous strains that still retain a sense of rhythm. Repeat Bar is by far the most rhythmic, creating a set pattern of bars that can be set to mutate over time, which would be great to use to emulate patterns similar to shuffles, etc. but in a generative vein.

The Fixed Pattern is a good place to diverge to the Pattern Editor, which is designed similar to the time-note matrices seen in many commercial programs on the market in 2004. One can import MIDI information, 'draw' note information in the Editor, and create MIDI information from live input sources, allowing for a great deal of flexibility when a user wishes to have a bit more control over the structure of a composition. By feeding to the Fixed Pattern voice, the user has the ability to introduce mutation and note probability percentages, adding some uncertainty to the set phrasing.

The two remaining voices, Follows and Listening, have special functions and are dependent on other forms of input for their functioning. For example, the Follows voice imitates another voice, but with a variation in interval, time delay, or probability of repetition.



tion that allows for rounds, choral structures, and the like. The Listening voice allows the introduction of external MIDI data to be taken into the *Koan* environment live from any external MIDI device. In this way, *Koan* can be used -- since it also outputs MIDI -- as a generative MIDI event processor as well as a musical tool, among other possibilities.

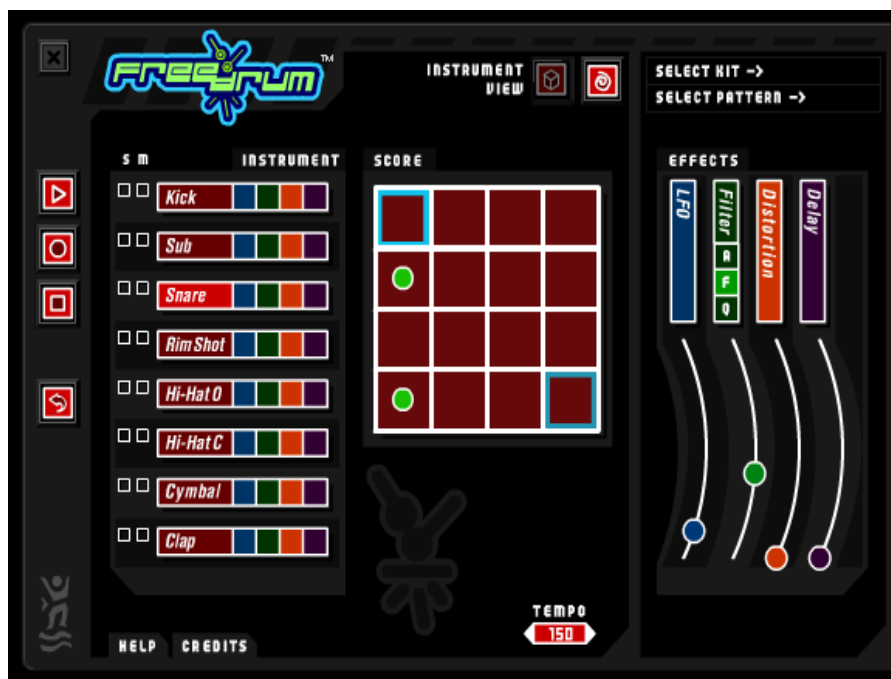
The third panel consists of the *Sseyo Koan Synth Engine*, which is quite unique to this application. The structure of the SKSE is similar to traditional modular synthesis in allowing the creation of synth modules that can be wired together using FLO's filters, particle systems, etc. The SKSE is actually a fully-featured modular soft synth that just happens to be part of the *Koan* system, allowing samples to be combined with envelopes from particle systems to create granular synthesis as well as the Buchla-esque patching-together of oscillators and filters to get some incredible retro sounds. Again, it's not overly pretty, but once you get the knack of the interface, the power of the system becomes evident.

As you may or may have not ascertained, the *Koan* music environment is a product that is in part a labor of love, but also a development tool for the creation of online content for Web and mobile applications. This was what took Sseyo's into its relationship with Tao and into the creation of its next-generation intent Sound System for mobile applications.

However, still contained within the *Koan* system is a full functionality for the creation of online content through audio 'vectors' for embedding within web pages. These small files, which are accessed through the *Koan* browser plugin, allow for the embedding of generative sound within the web environment. In addition, *Koan* also allows for full system level access to its API through Java, Javascript, and C++ objects on both the PC and Mac platforms. Due to this functionality, *Koan* was making content for clients in the UK in the 90s, passing on variables to Flash animations through Javascript and creating generative audiovisual pieces on the Net.

To go further into the possible applications within web and new media installations is outside the scope of this already very technical review, but it's my hope that you got an idea of the multifaceted possibilities of this tool. Although my assessment of *Koan* may verge on gushing its praises, I want to be clear about a few things. I reiterate that the tool itself is not pretty, but one of those with almost Baroque functionality that comes with an endless array of levers and gauges. It is a development tool, not a prosumer app. I had discussions with programmer Pete Cole about the proprietary nature of the *Koan* Music engine's equations, but this is a minor point compared to the number of parameters that the music engine does allow you to access. Due to Sseyo's concentration on the Tao intent Sound System, *Koan* has not had a major revision in over three years. All this being said, *Koan* is also one of the more overlooked options in the creation of generative music (or event processing) software. It's relatively inexpensive (\$199), has a great deal of flexibility and depth, and allows the creation of content across a wide variety of electronic media. It's not as flexible an environment as Max/MSP, but it's not intended for the same purposes either. Although I may make mention of some of *Koan*'s less attractive points such as its interface, I would metaphorically put them in the category of Cindy Crawford's mole, which is considered a beauty mark.

Ranking 4/5



*FreeDrum, a Koan-powered
Flash Applet by Andrew
Garton /ToySatellite.org*